Tapping with a Handheld Stick in VR: Redirection Detection Thresholds for Passive Haptic Feedback

Yuqi Zhou*
Purdue University

Voicu Popescu†
Purdue University

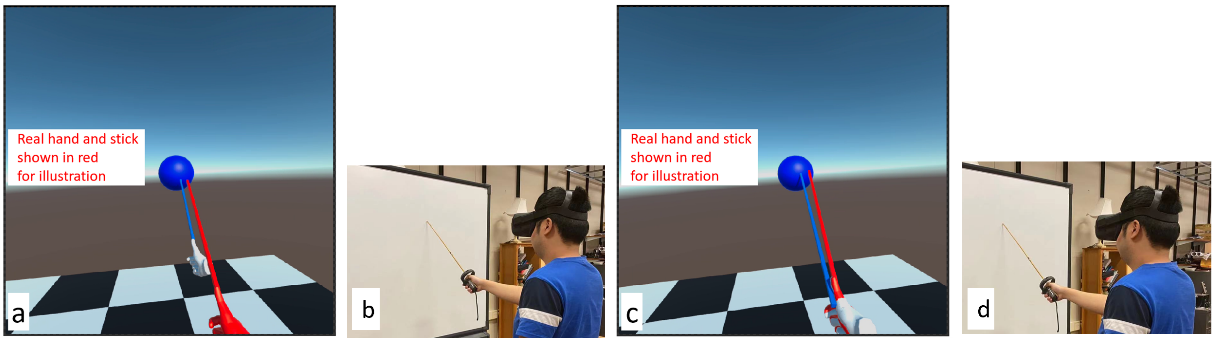

Figure 1: Passive haptic redirection methods for tapping with a handheld stick in VR. For *DriftingHand* (left), the virtual hand is translated forward to bridge the gap between the virtual object (blue sphere in VR image a) and real object (whiteboard in video frame b); the virtual stick is parallel to and of the same length as the real stick. For *VariStick* (right), the virtual hand stays in place (c), and the real to virtual gap is bridged by scaling the virtual stick. The real hand and stick are shown in red for illustration purposes.

Abstract

This paper investigates providing grounded passive haptic feedback to a user of a VR application through a handheld stick with which the user taps virtual objects. Such an investigation benefits VR applications beyond those where the stick interaction is actually an integral part of the narrative. Providing passive haptic feedback through a handheld stick as opposed to directly through the user's hand has the potential for more believable and more frequent feedback opportunities. The stick is likely to dull the user's haptics perception and proprioception, potentially avoiding a haptics perception uncanny valley and increasing the redirection detection thresholds. Two haptics redirection methods are proposed: the *DriftingHand* method, which alters the position of the user's virtual hand, and the *VariStick* method, which alters the length of the virtual stick. Detection thresholds were measured in a user study (N = 60) by testing the two methods for a range of offsets between the virtual and the real object, for multiple stick lengths, and multiple distances from the user to the real object. Overall, the study reveals that *VariStick* and *DriftingHand* provide an undetectable range of offsets of [-20cm, +13cm] and [-11cm, +11cm], respectively.

Keywords: Grounded passive haptics redirection, haptic retargeting, detection thresholds, handheld prop, virtual reality.

Index Terms: Human-centered computing—Visualization—Visualization techniques—Treemaps; Human-centered computing—Visualization—Visualization design and evaluation methods

1 Introduction

Current virtual reality (VR) technology provides the user with vivid images of complex virtual environments (VEs), tailored to the current user's view with little latency. All-in-one VR headsets now do on-board rendering and inside-looking-out tracking, at consumer-level price points, without the need of external graphics workstations or trackers. However, the illusion of actually being present in the VE requires catering to more than the user's vision, and VR technology progress towards fooling all of the user's senses has been harder to come by. In particular, providing the user with haptic feedback as they make contact with virtual objects is both important and challenging, as noted by many VR haptics researchers [48].

Gloves [12], suits [32], pens [35], and vibrating hand controllers [3] are active haptics approaches that work well for applications where evoking a low intensity haptic feedback is sufficient. Grounded passive haptics approaches rely on objects of the real world hosting the VR application to give the user the illusion of making contact with objects of the VE. For example, the user can sit in a virtual chair if a real-world chair exists with the same position, orientation, and shape as the virtual chair. However, the real world does not match the VE one to one, and passive haptics opportunities are scarce [16]. This scarcity has been addressed through *redirection*, a general VR interaction paradigm that can create more passive haptics opportunities by fudging discrepancies between the real and the virtual. For example, a user's hand can be redirected to grab the same real cup although they reach for different virtual cups [6], or a user's virtual index finger can be redirected to make contact with a button of a virtual airplane cockpit as their real finger touches a concave real-world plywood prop placed in front of the user [16].

In this paper we investigate providing a user of a VR application with grounded passive haptic feedback through a handheld stick: the user holds a handheld controller extended with a real-world bamboo stick, the stick is tracked and rendered in the VE, and the user taps virtual objects with the virtual stick. We investigate providing haptic feedback through the handheld stick rather than directly to the user's hand for two reasons. First, there are VR applications where the handheld stick VR interaction is an actual part of the application's narrative, e.g., to fight a virtual enemy [11], to pop virtual balloons, or to play a drum set in VR [29]. Second, compared to providing haptic feedback directly to the user, the stick provides an additional level of indirection, which we hypothesize could be exploited by *all* VR applications to provide haptic feedback more frequently.

*e-mail: zhou1168@purdue.edu
†e-mail: popescu@purdue.edu

Indeed, it is reasonable to expect that the acuity of the user's haptic feedback perception diminishes when tapping a real object with a stick compared to actually touching the object with their own hand. Touching with one's hand provides high resolution multidimensional haptic feedback, e.g. a precise indication of a small object's shape, of a wall's normal, or of whether a surface is wet, smooth, or cold. When the haptic feedback fails to deliver these expected sensations, the user is likely to perceive the discrepancy as salient, surprising, and disappointing, which could reduce their sense of presence in the VE. Conversely, when the user touches an object with a stick, if the real and virtual taps are well synchronized, and if visual contact cues are provided, the user is likely to find the interaction believable, since the user's visual sense is dominant [6, 15]. The stick adapts what the user can perceive to what passive haptics can provide, which holds the potential to avoid a haptic perception "uncanny valley" [9].

Furthermore, the stick appends a foreign segment to the user's arm that the user's proprioception is less familiar with and therefore less likely to keep track of accurately, compared to the user's hand. Consequently, we hypothesize that tapping with a stick should increase the amount of redirection that can be used by the VR application without it becoming objectionable. This hypothesis is anchored by prior work on proprioceptive stimulation [45, 46] and on versatile VR handheld props [40, 53] In other words, tapping with a stick *should* allow for larger discrepancies between the position of the virtual object being tapped and the position of the real-world object providing the passive feedback. The wider range of redirection parameter values will translate to more frequent opportunities for passive haptic feedback, as more virtual objects can be aligned with a given set of static real-world objects.

We have devised and investigated two methods for redirection when providing grounded passive haptic feedback through a handheld stick in VR. A first method, dubbed *DriftingHand*, alters the position of the user's virtual hand, and thereby that of the tip of the virtual stick, such that the virtual tip touches the virtual object when the tip of the real stick touches the real object (see Fig. 1, left). The *DriftingHand* approach does not modify the length of the virtual stick. A second method, dubbed *VariStick*, alters the virtual stick length and orientation such that the virtual and real tapping are synchronized (see Fig. 1, right). The *VariStick* method does not alter the position of the end of the virtual stick held by the user, i.e., it does not alter the position of the user's virtual hand.

We have conducted a user study where (N = 60) participants tapped virtual spheres with a handheld stick. Grounded haptic feedback was provided by a real whiteboard in front of the participants. The virtual spheres were placed at various distances behind or in front of the whiteboard, and the tapping actions in the real and virtual world were aligned using the *DriftingHand* and the *VariStick* methods. The study measured the detection threshold for the two methods using a two-interval forced choice (2IFC) design. 2IFC is a variant of the general two-alternatives forced choice (2AFC) design where the participant is not exposed to the two stimuli simultaneously, but rather in succession. A participant taps two spheres, one at the time, for one of the spheres haptic redirection is applied, and the participant is asked to choose the sphere where redirection was applied, whether they think they detect the redirection or not. The detection threshold is the largest distance between the real and the virtual object for which users cannot detect the redirection reliably. The study measured the detection threshold separately for when the sphere was behind or in front the whiteboard. The study estimated a total of 48 detection thresholds: 2 thresholds for each of 2 methods for each of 4 stick lengths and for each of 3 user to whiteboard distances. The detection thresholds were in the [-20cm, +13cm] range for *VariStick* and in the [-11cm, +11cm] range for *DriftingHand*.

We also refer the reader to the video accompanying this paper, which illustrates the two passive haptic redirection methods investigated by this paper, as well as the study in which they were evaluated.

2 Prior Work

Haptic feedback allows users to not just see but also feel virtual objects. Like others [38], we partition the discussion of prior work on haptics into active and passive methods.

2.1 Active and Mixed Haptics

In active haptic feedback methods, the VR system triggers actuators in contact with the user to convey haptic feedback. a plethora of VR haptic gloves [2, 12, 24] and suits [1] have been developed. The actuators are attached to the user who takes them along when moving, which affords a large range of motion. However, wearable haptics can be intrusive, as the actuators are always in contact with the user, and not just when they are needed to provide haptic feedback. Furthermore, the haptic feedback is of limited spatial and intensity resolution. Researchers have increased simulation fidelity for haptic feedback to evoke more closely specific different real-world surfaces. For example, Benko et al. developed the NormalTouch handheld controller, which renders haptic feedback along different directions using a tiltable and ex-trudable platform, and the TextureTouch handheld controller, which renders haptic feedback corresponding to object shape and surface structure using a matrix of 4x4 pins [8].

A class of haptics approaches rely on user-held devices, or props. With the classic PHANTOM haptic system, the user holds a pen-like controller with which they explore virtual surfaces [35]. The pen is tracked with a five or six degrees of freedom mechanical arm that also provides haptic feedback to the user by applying forces to the pen. High frequency bursts allow simulating surface bumps and rugosity. The PHANTOM was used in the Nanomanipulator project as an interface to an Atomic Force Microscope [27]. The PHANTOM is not portable, so the user is confined to a seated exploration of the VE. The PHANTOM controller is an example of a prop that the user relies on to experience the VE. Another example is a handheld VR controller which vibrates [3], for example, when a user hits a ball in a VR tennis game.

One limitation of props used for haptic feedback is their lack of versatility. For example, a plastic sword with a thin, flexible blade might be just what a foil-fencing VR app needs, but the same prop cannot replicate the feel of a hammer, as needed by a second application. Zenner and Kruger [53] investigated a versatile prop in form a stick which can redistribute its weight dynamically to simulate different handheld objects as needed by the application. The versatile prop was later used in conjunction with retargeting to further increase the range of scenarios where it could provide haptic feedback convincingly [56].

Props can serve as a natural interface between the real and the virtual. A virtual replica of the prop is rendered in the VE, the user sees the virtual prop, and the virtual prop is well integrated into the narrative of the VR application. Therefore the user is aware of the prop and manipulates it comfortably and adroitly based on real-world experience, as opposed to an ambiguous haptic glove, which the user feels all the time (in the real world), but never sees (in the virtual world). Hoping to capitalize on these advantages of a handheld prop, we explore providing haptic feedback to the user through a handheld stick. However, we investigate a passive and not an active haptic feedback approach.

An alternative is to do away with the actuators or props in direct contact with the user, and to rely on the VR system to actively modifies the physical world. Such hybrid active/passive haptic systems are called *encountered-type haptic displays* (ETHDs). An example of ETHD is the Robotic Shape Display, which uses a robot to place a physical board at locations aligned with virtual objects with which the user interacts [36]. Recently, five interaction techniques designed for ETHDs have been studied in the context of a virtual drawing task [37], to shed light on how to best address common issues with ETHDs, such as small contact areas, latency, and inadvertent collisions with the user. Fundamental challenges of ETHDs include limited range and high equipment cost.

2.2 Passive Haptics

In passive haptics, the real-world object providing the haptic feedback is stationary, and, in order for the user to feel like they make contact with a virtual object, the real and virtual objects have to be aligned. The advantages are simplicity, avoiding the user encumbrance specific to haptic gloves and suits, as well as the realism of interacting with a real-world object.

Passive haptic feedback works well when there is a one-to-one mapping between a real world and a VE object. For example, an early touchscreen like interface was developed by giving the user of a VR application a wooden paddle replicated in the virtual world, that could be used to provide 2D input with the index finger [33, 34]. Another example is the use a physical model of the brain in neurosurgical training [28].

The main challenge in passive haptics is the scarcity of haptic feedback opportunities, as haptic feedback can only be conveyed to the user when the virtual target happens to be aligned with a real-world object. With a one-to-one mapping between the real and the virtual world, providing passive haptic feedback for the entire VE requires the availability of a high-fidelity physical replica of the VE. Instead of creating the physical environment to match the VE, it is easier to constrain the VE to match the general layout of available physical spaces, as in substitutional reality [20]. A second approach is to design physical objects that can be used to cover the haptic feedback needs of many, or even *all*, possible virtual objects [17].

A third approach for increasing the availability of passive haptic feedback opportunities is redirection. Passive haptic redirection is an adaptation of the redirected walking (RDW) approach developed for achieving VR navigation that is robust to real/virtual world size mismatches [39]. Examples of using redirection in passive haptics are *redirected touching* [31], and *haptic retargeting* [6], which map multiple virtual objects to the same real-world object through a combination of user arm and VE geometry warping. Another example is the use of *sparse haptic proxies* [16], where a real-world object that approximates a piece of complex VE geometry is used to provide haptic feedback, such as a concave hemispherical prop that provides haptic feedback when the user pushes the buttons of a virtual cockpit. A research testbed for redirected touching is now available [52]. Other than static objects, redirection was also applied to physical proxies with movable parts [23].

An important sub-problem for all VR redirection approaches is to determine the thresholds below which the manipulation goes unnoticed by users [7, 10, 14, 21, 25, 26, 42, 43]. In RDW, researchers investigated how much user rotation and translation can be modified [26], and, conversely, how much space is needed to achieve ”endless” navigation in large VEs [5]. We discuss the prior work on the *methodology* of detection threshold estimation in Sect. 4.1.

Compared to RDW, the detection thresholds in passive haptic redirection have been studied less. The *redirected touching* [31] and the *haptic retargeting* [6] works do not estimate detection thresholds. Detection thresholds have been estimated for hand redirection in tasks where the user is asked to judge the location of their redirected virtual hand based on where they see it in the virtual world [54]. Whereas these findings are a first step towards determining the detection threshold for passive haptic feedback redirection, actually touching a physical object provides substantially more information to the user, and the thresholds are likely to be different.

Researchers have also estimated detection thresholds for *ungrounded* haptic retargeting, where the user holds the real object providing feedback. For example, a physical grabbing tool like a pair of chopsticks can provide haptic feedback by fully closing even if the corresponding virtual tool is retargeted to close partially to grab a virtual object [51]; the study finds that the travel distance difference between the virtual and real chopsticks has to be in the [-1.48cm, 1.95cm] range. In another study [44], physical tools like a hammer and a screwdriver have their VR use extended using redirection, collecting subjective user feedback, without estimating detection thresholds. Visually impaired VR users have been provided navigational support through a haptic cane [57]. Since the haptic feedback is provided by the tool, like in the case of a vibrating handheld controller, the haptic feedback is ungrounded. Researchers have also studied *grounded* passive haptic redirection, where the real object providing the haptic feedback is not held by the user but rather fixed in the real world scene. One such study took advantage of user blinks to hide the haptic redirection of the user's hand [55]. A challenge with establishing detection thresholds is the bias of self-attribution, which makes the user less likely to detect a redirection when it helps them complete the task [18, 19]. Furthermore, redirection detection thresholds should be measured along multiple directions, as each direction has its own visual implications, i.e., lateral changes are more visually salient than changes along the view direction [15, 47].

We build on this prior work on redirection for passive haptic feedback in VR. We investigate the detection thresholds when applying redirection to tapping with a handheld stick of various lengths, which has the potential to provide a natural interface between the virtual and the real world, to extend the reach of the user, and to be less familiar to the user's proprioception. Furthermore, we investigate how the distance from the user to the real world object affects the detection thresholds. Our research is anchored by psychology research results that have shown that the acuity with which users perceive objects decreases as the distance to the object increases [13]. We focus on redirection aligned with the view direction, as the stick is moved forward, and we leverage self-attribution, as the redirection facilitates the user's goal of hitting the wall.

3 Redirection for Passive Haptic Feedback when Tapping with a Handheld Stick in VR

When a user of a VR application makes contact with a virtual object, passive haptic feedback allows the user to perceive the contact through their sense of touch with the help of a stationary object of the real-world environment hosting the VR application; the user thinks they are making contact with the virtual object, whereas, in reality, they are making contact with the real object that provides the haptic feedback. For the illusion to work, the contact with the virtual and real objects have to be synchronized, which limits the opportunities for passive haptic feedback to the rare instances when the geometry of the real and virtual environments match. More opportunities for passive haptic feedback can be created if the VE and the user's position and orientation in the VE are altered through redirection to achieve a temporary alignment between the virtual object with which the user is about to make contact, and a real object that can provide the haptic feedback.

When the user makes contact with the virtual object directly with their hand, or their fist grasping the controller, it is difficult to redirect without the user noticing. One option is to move the virtual object as the user reaches for it such that the virtual object aligns with a real-world object at the assumed point of contact. Another option is to move the user's virtual hand such that the virtual contact is delayed or expedited to coincide with the real contact. For both approaches, the altering has to be diluted over a sufficiently long time interval that might not always be available in complex VEs where it is hard to anticipate the virtual object with which the user is going to want to make contact.

However, if the user relies on a handheld stick rather than their hand to probe the virtual world, redirection for passive haptic feedback might be easier to hide. Consider the scenario where the user holds a VR handheld controller to which a real stick of about 70cm has been attached as shown in Fig. 1. The stick is tracked and rendered in the VE. The stick serves as a connection between the user and the virtual object they intend to tap, and it presents the

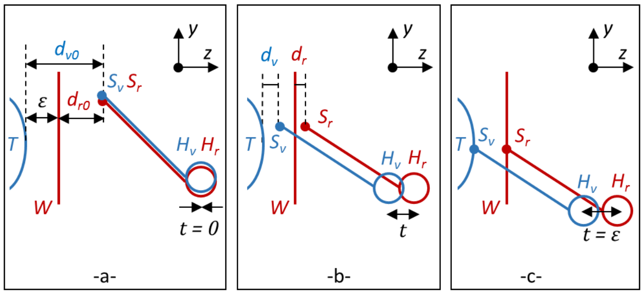

Figure 2: *DriftingHand* redirection shown in a vertical plane. Real/virtual entities are shown with red/blue. The user taps a virtual object T with a stick with tip S held in their hand H. Passive haptic feedback is provided by a wall W. When the stick tip gets sufficiently close to T, the redirection of the virtual hand H_v begins (a). Here T is farther than W, so H_v moves faster than the user's actual hand H_r to cover the additional distance ε (b). The real and virtual sticks hit T and W at the same time (c).

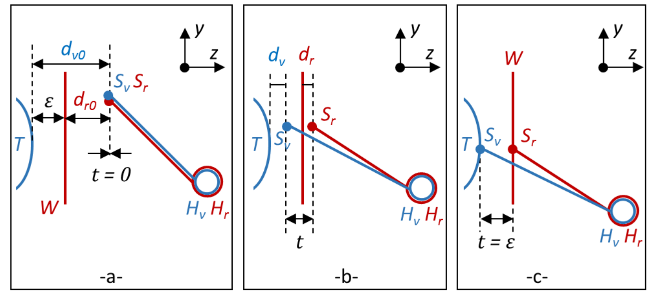

Figure 3: *VariStick* redirection shown in a vertical plane. Real/virtual entities are shown with red/blue. The user taps a virtual object T with a stick with tip S held in their hand H. Passive haptic feedback is provided by a wall W. When the stick tip gets sufficiently close to T, redirection begins (a). Redirection changes the length and orientation of the stick while keeping the virtual H_v and real H_r hands aligned (b). The real and virtual sticks hit T and W at the same time (c).

opportunity of a redirection approach where the altering focuses on the stick as opposed to on the virtual object or on the user. Such redirection might allow bridging larger distances between virtual and real objects to provide passive haptic feedback without the user noticing. Indeed, the user's mental tracking of the stick is likely to be less accurate than that of their hand. Furthermore, as the user approaches a virtual object to make contact, the stick shifts the user's focus from their hand to the tip of the stick, which could help hide the altering of the position of the user's virtual hand. Based on these assumptions, we have developed two passive haptic feedback methods designed specifically for the handheld stick VR interaction scenario, one that redirects the stick together with the user's hand, and another that concentrates the manipulation at the tip of the stick, without any change of the user's hand pose.

3.1 The *DriftingHand* redirection method

One redirection method, which we dubbed *DriftingHand*, alters the position of the virtual hand of the user such that the tip of the virtual stick makes contact with the virtual object at the same time the tip of the real stick makes contact with a real object. The method is illustrated in Fig. 2. The method keeps the virtual stick parallel to the real stick, and it does not change the length of the stick; the redirection is applied as a translation of the the virtual hand of the user and of the virtual stick, forward (here) or backward, i.e., along the z axis, to increase or decrease the user's virtual reach when the virtual object is behind or in front of the real object.

Beginning of redirection. When the tip of the stick gets within a horizontal distance d_{v0} of the virtual object T, and there is a real-world object W within a distance ε_{max} of T, redirection begins (panel a in Sect. 3.1). The distance ε between virtual and real contact objects can be negative, i.e., T is farther from the user than W, like in Sect. 3.1, or positive, i.e., T is closer to the user than W. If there is no real-world object sufficiently close, there is no redirection, and the virtual contact occurs without passive haptic feedback.

During redirection. While the distance d_r between S_r and W is less than the initial distance d_{r0} when redirection began, given by Equation 1, the virtual hand is offset from the real hand through a redirection translation t on the z axis given by Equation 2 (panel b).

$$d_{r0} = d_{v0} - \varepsilon \tag{1}$$

$$t = \varepsilon(1 - d_r/d_{r0}) \tag{2}$$

d_r is computed from the precalibrated position of the real-world object W, from the H_r provided by the handheld controller, and from the known stick length and orientation with respect to the controller. Initially, when $d_r = d_{r0}$, $t = 0$. Then, as d_r decreases to 0, t increases to ε. Similarly, the distance d_v between S_v and T is given by Equation 3, and it decreases from d_{v0} to 0.

$$d_v = \varepsilon + d_r - t \tag{3}$$

Contact. When S_r reaches W, both d_r and d_v are 0, so the real and virtual contacts occur simultaneously (panel c).

The approach has two parameters. One is ε_{max}, which is the maximum distance between the virtual and real objects that redirection attempts to bridge. We have investigated how large ε_{max} can be while redirection remains inconspicuous to users, as detailed in Sect. 4. The other parameter is d_{r0}, which is the distance between the tip of the real stick and the real wall when redirection begins. In practice we use $d_{r0} = 100cm$. ε is known based on the known real and virtual object positions, and d_{v0} is computed using Equation 1.

The *DriftingHand* method shifts the user's focus from their hand to the tip of the stick, so the user is less likely to see their virtual hand on which the redirection acts. However, the position of the tip of the stick is rigidly coupled to that of the virtual hand, so it inherits the position alterations applied to the virtual hand, which are in contradiction with the user's proprioception.

3.2 The *VariStick* redirection method

A second redirection method, which we dubbed *VariStick*, alters the length of the virtual stick and its orientation such that the tip of the virtual stick makes contact with the virtual object at the same time the tip of the real stick makes contact with a real object. The method is illustrated in Fig. 3. The tip of the virtual stick S_v is translated forward to reach T which is farther from the user than W. The virtual hand H_v of the user is kept in place, which requires scaling up the virtual stick to increase its length. Similarly, when the virtual object is closer than the real object, S_v lags behind S_r, and the virtual stick is scaled down to reduce its length.

The translation t of the tip of the virtual stick is exactly the same as the translation of the virtual hand described for the *DriftingHand* method (see Equations 1 and 2). The translation starts out at 0 when the redirection begins, and then it increases to ε as redirection proceeds, to bridge the gap between T and W. Given S_v, the length of the stick is computed as the distance between the user's hand and S_v, which puts the tip of the virtual stick at the desired location while keeping the virtual hand in place.

Like the *DriftingHand* method, the *VariStick* method also uses the stick to shift the user's focus away from their hand. Furthermore, *VariStick* also alters the position of the tip of the virtual stick to achieve virtual contact without moving the virtual object. However, unlike *DriftingHand*, *VariStick* keeps the virtual hand in place. The redirection is introduced gradually, from the grip to the tip of the stick, through scaling, potentially avoiding a proprioception conflict.

4 EVALUATION

Our goal is to evaluate the *DriftingHand* and the *VariStick* redirection methods for haptic feedback in terms of the range of virtual to real object distances ε that each method can bridge without the user noticing. We first describe the experimental design framework (Sect. 4.1), and then the user study (Sect. 4.2).

4.1 Experimental design framework

Psycho-physics research has established over a century and a half ago that detection thresholds should be measured with a forced choice experimental design [22]. In such a design the participant is presented with two stimulus levels, one of which could be 0, and is asked to choose the higher stimulus level. The participant does not have the option of saying that they are not sure, or don't know, i.e., they are forced to select one of the two. One variant of the two-alternative forced choice (2AFC) design presents to the user both alternatives at the same time. In VR research, one cannot let the user experience both alternatives simultaneously. Consequently the alternatives are presented in succession, and the forced-choice occurs over a time interval, a variant of the 2AFC experimental design that is called two-interval forced choice design (2IFC) [50].

When the participant does not detect the level, or cannot discern between the two levels, the accurate response rate is 50%, corresponding to chance behavior. The 50% rate is the point of subjective equivalence (PSE). An important goal of psychometric research, which we share, is to establish the threshold for the stimulus level above which a participant is expected to choose the correct alternative consistently. Determining the threshold is usually done by fitting a sigmoid psychometric function to the correct response rates over various stimuli levels, and then intersecting the sigmoid with the 75% line to obtain the threshold.

In redirected walking (RDW) VR research, establishing the rotation and translation gain detection thresholds with a design based on 2AFC tasks requires that the participant walk in the virtual environment too many times for the design to be practical. Instead, RDW detection threshold research relies on a simplified version of the 2AFC design, where the participant walks the VE only once and is then asked which way the path bent, left, or right [43]. In other words, the participant is not presented with two levels of the stimulus, but rather just one. Such simplified, two-alternative one-stimulus-level task has been called a pseudo-2AFC design [26], or a "yes/no" design [30]. While they present a practical way of estimating thresholds, these single stimulus designs can be affected by bias when participants consistently choose one answer when they do not know the correct answer [50].

For the problem at hand, i.e., determining the detection thresholds in redirected passive haptic feedback, using a rigorous 2AFC experimental design is tractable, as the user can hit a virtual object quickly and effortlessly, so experiments with many trials are practical.

Another aspect of the experimental design is whether to use fixed stimuli levels, or whether to adapt the stimulus level based on the participant's previous responses [30]. The advantage of adaptive approaches is a reduced number of trials. The advantage of constant level approaches is that they do not only provide the threshold value, but also the shape of the detection rate as a function of stimulus level. Furthermore, constant level approaches are more robust, avoiding the perils of incorrect convergence that could affect adaptive methods. Since the number of trials is less of a concern in our case, we have chosen a constant level approach, where all participants are exposed to the same set of predetermined stimulus levels.

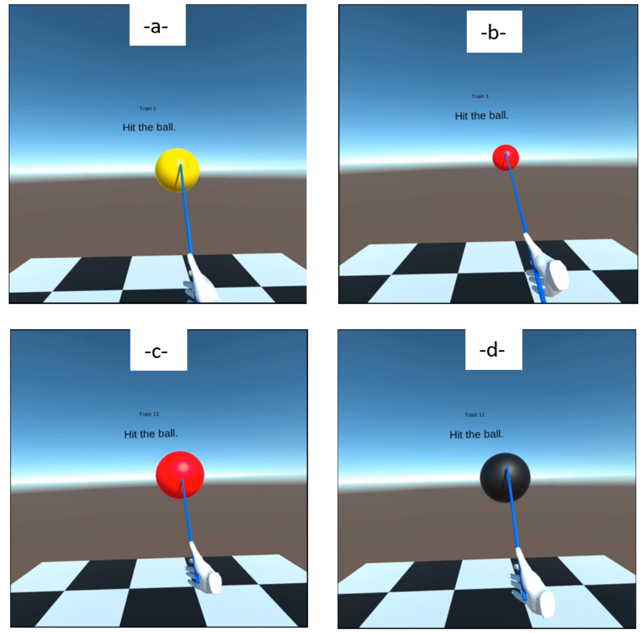

Figure 4: User study frames. Each row shows the two spheres of the same trial. The participant sees the spheres one at the time.

4.2 User study

We investigated the distance offset detection thresholds for the two passive haptic redirection methods in a user study (N = 60). The study was approved by our university's Institutional Review Board.

Participants. We have recruited 60 participants with ages between 19 and 30 from the undergraduate and graduate student population of our university. 15 were female. 12 participants self-reported that they had no experience with VR, 15 used VR once, 29 used VR occasionally, and 4 used VR frequently.

Experimental setup. Each participant was asked to stand in front of a whiteboard (Fig. 1 b and c), wearing the VR headset, and holding the controller with a bamboo stick attached. The stick is sufficiently rigid to not bend under gravity, but it is also somewhat elastic to bend (and not break) when tapping an object. The VE shows the virtual hand of the participant holding a virtual stick (Fig. 1 a and c, but without the real hand and stick visualizations which were added to the figure for illustration purpose). We implemented our methods in Unity 3D (version 2020.2.7f1, [4]) and the VR application ran on an Oculus Quest headset [3].

Procedure. We used a two-interval forced choice (2IFC) design. For each trial, the participant was shown two virtual spheres, one at the time. The spheres appeared at the height of the participant's head and in front of the participant's hand holding the stick. This left/right hand agnostic design allows the participant to hold the stick with their dominant hand. Once a sphere appeared, the position of the sphere remained fixed, i.e., it did not change as the participant moved their head or their hand. The depth of one sphere was always chosen such that the sphere be behind the whiteboard and tangent to the whiteboard plane, i.e., there was no distance between the virtual object the user hit, and the real object providing the feedback. The depth of the other sphere was offset with respect to the whiteboard plane, i.e., our haptic redirection methods had to bridge this offset to synchronize the contact between the virtual stick and the virtual sphere with the contact between the real stick and the whiteboard. The order in which the two spheres appeared was randomized. For each trial, the offset ε of the sphere on which redirection was applied was counterbalanced from the set {-100cm, -75cm, -50cm, -25cm, -10cm, 0cm, 10cm, 20cm, 30cm}. Negative ε values correspond to sphere positions behind the whiteboard, and positive ε values correspond to sphere positions in front of the whiteboard. For the other sphere $\varepsilon = 0cm$. The trials were grouped in blocks of 27 trials, with each of the 9 epsilon values being used exactly three times.

For each trial, the order in which the two spheres appeared was randomized. The sphere colors changed randomly from trial to trial from the set {green, red, yellow, and black}, and the two spheres of a trial always had different colors. The spheres also changed diameter randomly from trial to trial in the interval [25cm, 35cm]. The change in sphere color and size aims to prevent participants from learning which sphere is aligned with the whiteboard. In Fig. 4, the spheres aligned with the whiteboard are a, d.

The participant was asked to hit each sphere with the stick. When contact was made, the sphere moved back, away from the user, with a velocity proportional to the speed of the stick (see accompanying video). Contact was also illustrated visually by the shadow cast on the sphere by the tip of the virtual stick, which touched the tip of the virtual stick on contact. Participants received haptic feedback from the whiteboard whose collision with the real stick was synchronized with the collision between the virtual stick and the virtual sphere using our two redirection methods.

After the participant hit the second sphere, the VE displayed a billboard with the question "For which sphere did the virtual stick change in length?" for the *VariStick* method, and "For which sphere did your virtual hand NOT move as expected?" for the *DriftingHand* method. The billboard provided two choices, identifying the spheres by their color and the order in which they appeared, e.g., "1. Black" to the left and "2. Yellow" to the right for a trial where a black sphere appeared first, followed by a yellow sphere. The participant selected an answer with the virtual laser method. The participant could not advance to the next trial without selecting one of the two choices. The participant was given the option of redoing a trial exactly the same way as the first time using a "Redo" button that appeared on the question billboard.

Dependent variables. Each block of trials measured two detection thresholds, in cm: τ_f, which is the detection threshold for the distance by which the virtual object is farther from the user than the real object, i.e., the virtual object is behind the real object ($\varepsilon < 0$); and (2) τ_n, which is the detection threshold for the distance by which the virtual object is closer to the user than the real object ($\varepsilon > 0$), i.e., the virtual object is in front of the real object. We distinguish between τ_f and τ_n because it is likely that participants detect the redirection more easily when the virtual object is closer to them than when it is farther. A pilot study revealed that participants can easily detect the redirection for $\varepsilon = -100cm$ and $\varepsilon = 30cm$, so our range of ε trial values contains the two detection thresholds.

Independent variables. We measured τ_f and τ_n for 24 conditions: 2 haptic redirection methods (*VariStick* and *DriftingHand*) $\times$ 4 stick lengths L (30cm, 50cm, 70cm, 90cm) $\times$ 3 distances between the user and the whiteboard d_w (d_{w0} - 15cm, d_{w0}, and d_{w0} + 15cm), where d_{w0} is the default tapping distance as a function of the stick length. The different stick lengths were implemented with different bamboo sticks of the appropriate lengths. Based on preliminary experiments, the default tapping distances were set to 76cm, 88cm, 100cm, and 112cm, for stick lengths of 30cm, 50cm, 70cm, and 90cm, respectively. The whiteboard was on wheels and the distance between the user and the whiteboard was modulated by the experimenter moving the whiteboard front and back between blocks of trials using precalibrated floor markings. The participant who was wearing the HMD did not see the whiteboard move.

Data collection. In order to keep the total involvement time for a participant below 30 minutes, each participant used only one stick length, i.e., there were 15 participants assigned to each of the four stick lengths. Each participant started with a training session where they performed 18 trials, one for each of the 9 ε values, for each of the two redirection methods. During training, the participants received verbal instructions on what the questions mean, with examples of correct answers for the large ε values where the redirection was obvious. The participants were also told to take a one minute break any time their arm got tired. Then each participant performed 6 blocks of 27 trials, i.e., one block for each of the two methods and each of the 3 whiteboard distances. Each participant was involved in a single experimental session, which took at most 27min. During the session, the experimenter would check in with the participant periodically (approximately every five minutes) to ask how far along the participant was in their session. For each participant we collected the correctness of their 27 answers for each of their 6 blocks.

Data analysis. We estimated the detection thresholds τ_f and τ_n for each of the 24 condition combinations. For each combination, τ_f was estimated by fitting a psychometric function through six correct response rate points, i.e., one for each of the six $\varepsilon \leq 0$ values. The correct response rate for a given ε is computed over all 45 answers collected for that ε in that condition: 15 participants per condition times three repetitions for each ε. Similarly, τ_n was estimated from four correct response rate points, one for each $\varepsilon \geq 0$ values.

The psychometric function [49, 50] is given in Eq. 4, where it is defined based on the logistic sigmoid function shown in Eq. 5.

$$\Psi(\varepsilon;\alpha,\beta,\gamma,\lambda) = \gamma + (1-\gamma-\lambda)\Sigma(\varepsilon;\alpha,\beta) \tag{4}$$

$$\Sigma(\varepsilon;\alpha,\beta) = \frac{1}{1+e^{\delta\frac{\varepsilon-\alpha}{\beta}}} \tag{5}$$

Ψ is a function of the haptic redirection distance ε and returns the expected correct response rate. Ψ has five parameters. γ and δ define the overall shape of the psychometric function and are chosen based on the experiment. γ defines the lower range of the psychometric function; in all two forced choice experiments, like ours, the lower range of the correct response rate is 0.5, hence we set $\gamma = 0.5$. δ defines whether the psychometric function increases or decreases, so we use $\delta = +1$ for τ_f and $\delta = -1$ for τ_n.

The other three parameters (α,β,λ) modulate the shape of the psychometric function and they are fitted to the data points through an optimization that minimizes the square distance from the points to the curve. α controls the left/right translation. β controls the lateral scaling of the sigmoid, i.e., the range of ε values over which the correct response rate transitions between the chance and maximum levels. β is inversely related to the slope of the sigmoid over this transition region. λ fine-tunes the upper range to $1.0 - \lambda$ to account for the fact that data might not contain points with a perfect correct response rate of 1.0, since in some trials participants might forget the correct response or their concentration might lapse momentarily. We report the square root of the minimum, which has units like those of the y axis, i.e., the correct response rate. For example, a fitting error of 0.03 represents a 3% correct response rate interval.

Once the psychometric functions are fitted, the haptic redirection thresholds τ_f and τ_n are estimated by intersecting each psychometric function with the $y = 0.75$ line [49, 50].

Results. The 2x24 τ_f and τ_n detection threshold values are given in Table 1. They are also shown in graphical form in Fig. 6. The transition widths β of the fitted sigmoids are given in Table **??**. Fig. 5 shows the fitting of the psychometric functions for the 16 thresholds for $d_w = d_{w0}$. The fitting errors are small, and largest when the correct response rate for 0cm is farthest from 0.5.

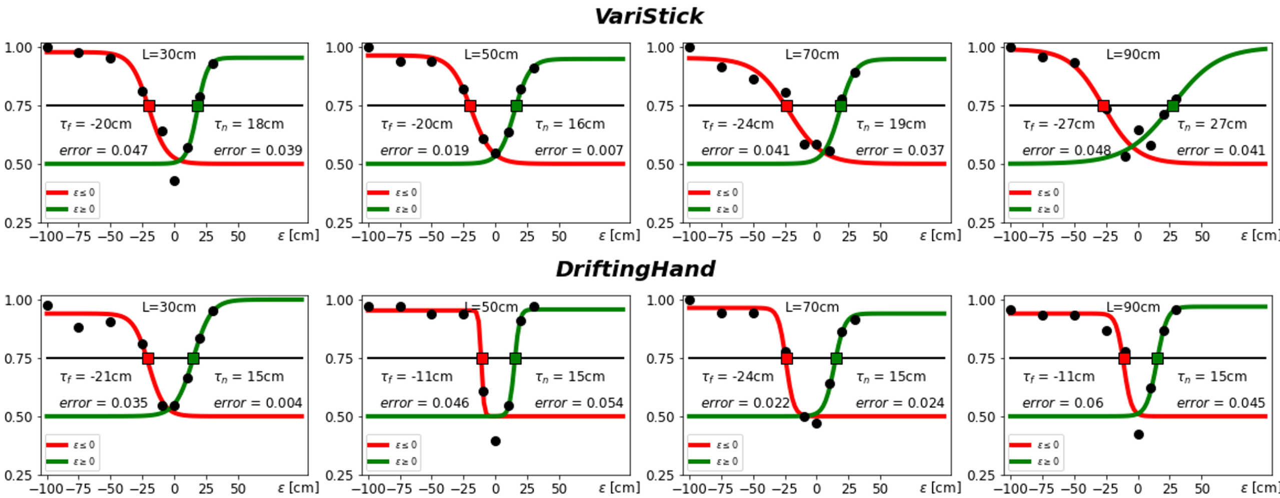

Figure 5: Psychometric function fitting for the two haptic redirection methods and for the four stick lengths L. The y axis gives the correct response rate. The detection thresholds τ_f and τ_n are estimated by intersecting the red and green curves with y = 0.75. For these graphs, the whiteboard was at the default distance d_{w0}; the graphs for the other two whiteboard distances are similar.

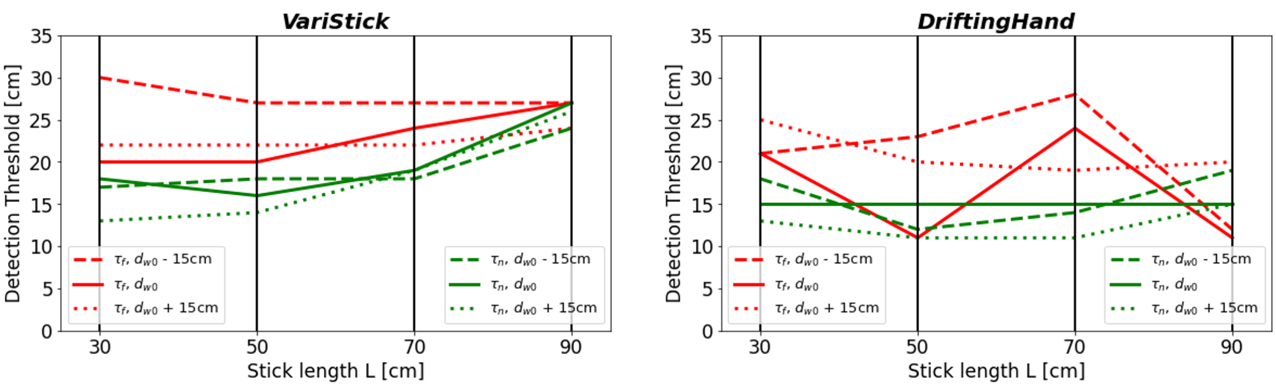

Figure 6: Detection thresholds τ_f ($\varepsilon \leq 0$) and τ_n ($\varepsilon \geq 0$) as a function of stick length L, for the two methods and three whiteboard distances.

		Stick length L [cm]							
		30		**50**		**70**		**90**	
Method	**d_w [cm]**	τ_f	τ_n	τ_f	τ_n	τ_f	τ_n	τ_f	τ_n
VariStick	d_{w0} + 15	22	13	22	14	22	19	24	26
	d_{w0}	20	18	20	16	24	19	27	27
	d_{w0} - 15	30	17	27	18	27	18	27	24
Drifting-Hand	d_{w0} + 15	25	13	20	11	19	11	20	15
	d_{w0}	21	15	11	15	24	15	11	15
	d_{w0} - 15	21	18	23	12	28	14	12	19

Table 1: Detection thresholds τ_f and τ_n in cm for the two methods, the four stick lengths L, and the three whiteboard distances d_w.

Discussion. For the *VariStick* method, the smallest τ_f and τ_n are 20cm and 13cm, which means that no matter the stick length and the whiteboard distance, participants were unlikely to detect redirections in the ε interval of [-20cm, 13cm]. For *DriftingHand* this interval was smaller, i.e., [-11cm, 11cm]. Fig. 6 shows that a piecewise linear curve for *VariStick* (left) is usually higher than the corresponding curve for *DriftingHand*. This suggests that overall *VariStick* hides the haptic redirection better.

For both methods τ_f tends to be larger than τ_n. One reason is that participants have much better visual distance judging ability in their proximity compared to farther away, which remains true in VR due to the stereo depth cues provided by the headset. Another reason is that, like in redirected walking, participants are more tolerant of superlinear gains than of sublinear gains: a participant is more likely to ignore the redirection when they see the stick move faster than expected than when they see it move slower. A third reason is that hitting a nearby object with a long stick is cumbersome and participants proceeded with great care, which was slowed down further by the sublinear gain, amounting to a slow motion tapping that allowed participants to discriminate more accurately.

For *VariStick*, the thresholds tend to increase with stick length, which confirms our hypothesis that the farther the point of contact is away from the end of the user's hand, the less the user can discern the redirection. The exception is for the shortest stick (L = 30cm) and the farthest whiteboard ($d_w = d_{w0}$ - 15cm), when participants had to poke the whiteboard barely reaching it with an almost horizontal stick, which had a tiny projection on the participant's view, hence the participant could not detect the change in length. Whereas this condition produced large thresholds, it is uncomfortable for the participant and unlikely to be useful in applications. For *Drifting-Hand*, the thresholds do not increase with distance, as the redirection is revealed to the user through the visualization of their virtual hand, and not at the far end of the stick.

As expected, the distance to the whiteboard, which was varied with a $\pm$15cm offset, is less of a factor for longer sticks, as illustrated by the fact that the piecewise linear curves converge as the stick length increases. This means that for longer sticks a VR application can use reliably the same detection threshold to provide passive haptic feedback to the user through real objects located at a range of distances relative to the user. Indeed, increasing the number of haptic feedback opportunities depends not only on the distance between the real and the virtual object that can be hidden from the user, but also on finding real objects at the appropriate distance to the user.

The main focus of this study is to estimate detection thresholds for each method, as described above. A secondary goal is to compare the near and far detection thresholds within and across methods. In order to investigate whether differences in detection thresholds are significant, one option is to estimate detection thresholds for each participant, and then to run a statistical test to compare the detection threshold means over all participants. However, this requires fitting the psychometric function to data points obtained by averaging three yes-no answers. This data is insufficient for a robust analysis.

A second option, which we adopt, is to run the statistical analysis on the thresholds estimated robustly for near and far, and for each method, over 12 conditions, i.e., 4 stick lengths × 3 wall distances. The average τ_n and τ_f for *VariStick* are 19.08±4.42 and 24.33±3.23; for *DriftingHand* they are 14.42±2.47 and 19.58±5.57. Table 2 shows a statistical comparison of these averages. We compare two averages using a paired t-test, as appropriate for our within-subjects design. We used the Shapiro-Wilk test [41] to investigate whether any of the differences has a non-normal distribution. The null hypothesis could not be rejected for any of the differences ($p > 0.05$). τ_f is significantly larger than τ_n within each method, which confirms our expectation that participants are less sensitive to the redirection when the target is farther away. Furthermore, *VariStick* has significantly larger detection thresholds compared to *DriftingHand*, for both near and far, which confirms our expectation that participants are more sensitive to redirections applied to the virtual hand.

VariStick		*DriftingHand*		*VariStick* vs *DriftingHand*			
τ_f-τ_n		τ_f-τ_n		τ_f-τ_f		τ_n-τ_n	
Norm.	t-test	Norm.	t-test	Norm.	t-test	Norm.	t-test
0.817	<0.001	0.221	0.012	0.198	0.012	0.590	0.001

Table 2: Statistical comparison of the differences between the thresholds τ_n and τ_f of the two methods. For each difference, the bottom row shows the p values for the normality and for the statistical test.

5 Conclusions, Limitations, and Future Work

We have devised and investigated two redirection methods for providing passive haptic feedback when the user of a VR application taps objects with a handheld stick. We have tested the two methods for four stick lengths and three user to real object distances. The cases when the virtual object was closer or farther than the real object were investigated separately. The study estimated a total of 48 detection thresholds.

Overall, the study reveals that the *VariStick* method buys the VR application more passive haptic opportunities by relaxing the alignment between the real and virtual objects to offsets in the [-20cm, +13cm] interval. This interval is only [-11cm, +11cm] for *DriftingHand*. The difference between the methods is expected as the user is likely to be less aware of the stick length and orientation than they are of the position of their hand.

The significance of the difference between near and far within a method, and the difference between methods has been confirmed with a paired t-test over the 12 thresholds estimated for each method and for near and far. Our study was designed to cover a large number of independent variables, i.e., stick length, wall distance, and method, which implies that each participant performed a limited number of trials, i.e., three, for each of the ε values and for each of combination of independent variables. Future work could extend the statistical comparison of the two methods on a participant by participant basis in studies where the number of repetitions is increased, at the detriment of fewer independent variables. Prior work measured hand redirection detection thresholds in purely visual VR tasks to be [-3.5cm, 6cm], i.e., 0.88-1.07 translation gain detection thresholds applied to a 50cm real hand translation [54]. The larger thresholds we found confirm our hypothesis that the stick extends detection thresholds, compared to the user's hand.

Like for any experiment based on an two-interval forced choice design (2IFC), which shows the two stimuli in succession and not simultaneously, one challenge is that participants might detect correctly the redirection but then forget for which sphere they detected it by the time the questions is displayed. We have tried to mitigate this problem by providing the redo button to participants. Another challenge is that, although short, the experiment is repetitive, which could disengage the participants and lead to mechanical answers. This implies that it is not easy to increase the number of trials in the hope of a more robust analysis. A possible solution is to increase the user's interest in the trials, for example by awarding points, or by playing back video and sound effects for correct answers, with the corresponding risk of introducing noise in the measurement.

In our study we used a vertical "plane" parallel to the user's image for the real-world object. Future studies might investigate different plane orientations, including single diagonal or horizontal planes, as well as multiple planes. Furthermore, we investigated detection thresholds for virtual to real offsets along the user view direction (z axis in Fig. 3), and future work should investigate offsets along the up-down (y) and left right (x) directions.

We have developed two methods for haptic redirection that are at opposite ends of the multidimensional space of haptic redirection design: one method avoids all manipulation of the user's hand position, whereas the other method manipulates it as needed to bridge the distance between virtual and real. Future work could investigate hybrid methods, where the two approaches are combined dynamically, e.g., more hand redirection when the hand is not visible, and more stick length change when the stick image footprint is small. Finally, whereas the redirections investigated were linear, future work could also examine changing the stick length or the hand position non-linearly, exploiting any non-linearity in the perception of such changes, e.g., faster change farther away from the user, or faster change at the beginning of the redirection while the distance between the stick and the target object is still large and the user is less likely to pay attention than when contact is imminent.

The present work takes a first step towards bringing the benefit of haptics to VR applications. Future work can rely on the measured thresholds to provide haptics in applications where the stick is integral to the narrative, for example in a VR fencing training application, where the user is battling the walls and furniture of the physical world. Future work should also explore the possibility for haptic feedback afforded by our method in the context of VR applications where the handheld stick is simply an interface between the user's hand and the virtual environment, like a virtual laser pointer with haptic feedback, assisting with fundamental VR tasks such as object selection. Finally, our findings can be applied to further extending the generality and versatility of VR props, such that they not only "feel right" in the user's hand by providing the proper weight distribution, but also feel right when used to act on virtual objects with convincing haptic feedback. Conversely, the weight redistribution of prior VR props could benefit our haptic redirection methods by helping hide the stick length changes and thereby increasing the detection thresholds.

Acknowledgments

We thank our study participants. We thank the reviewers for their feedback. We thank Purdue University's Computer Science Department for lending us the VR equipment used in this study.

REFERENCES

[1] bhaptics suit. https://www.bhaptics.com/. Accessed: 2020-05-15.

[2] Noitom hi5vr gloves. https://hi5vrglove.com/. Accessed: 2020-05-15.

[3] Oculus headset. https://www.oculus.com/. Accessed: 2020-05-15.

[4] Unity 3d engine. https://unity.com/. Accessed: 2020-05-15.

[5] M. Azmandian, T. Grechkin, M. T. Bolas, and E. A. Suma. Physical space requirements for redirected walking: How size and shape affect performance. In *ICAT-EGVE*, pp. 93–100, 2015.

[6] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*, pp. 1968–1979, 2016.

[7] B. Benda, S. Esmaeili, and E. D. Ragan. Determining detection thresholds for fixed positional offsets for virtual hand remapping in virtual reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 269–278. IEEE, 2020.

[8] H. Benko, C. Holz, M. Sinclair, and E. Ofek. Normaltouch and texturetouch: High-fidelity 3d haptic shape rendering on handheld virtual reality controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pp. 717–728, 2016.

[9] C. C. Berger, M. Gonzalez-Franco, E. Ofek, and K. Hinckley. The uncanny valley of haptics. *Science Robotics*, 3(17):Art–No, 2018.

[10] J. Bergström, A. Mottelson, and J. Knibbe. Resized grasping in vr: Estimating thresholds for object discrimination. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, pp. 1175–1183, 2019.

[11] A. Bieniek, A. Szczygioł, M. Chrzan, P. Wodarski, M. Morys, B. Bacik, G. Juras, R. Michnik, K. Paszek, and M. Gzik. Biomechatronic simulator for fencing training using virtual reality technology. In *International Conference Mechatronics*, pp. 30–37. Springer, 2017.

[12] M. Bouzit, G. Popescu, G. Burdea, and R. Boian. The rutgers master ii-nd force feedback glove. In *Proceedings 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002*, pp. 145–152. IEEE, 2002.

[13] J. R. Brockmole, C. C. Davoli, R. A. Abrams, and J. K. Witt. The world within reach: Effects of hand posture and tool use on visual cognition. *Current Directions in Psychological Science*, 22(1):38–44, 2013.

[14] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and F. P. Brooks. The hand is slower than the eye: A quantitative exploration of visual dominance over proprioception. In *IEEE Proceedings. VR 2005. Virtual Reality, 2005.*, pp. 3–10. IEEE, 2005.

[15] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and F. P. Brooks Jr. The hand is more easily fooled than the eye: Users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy. *Presence: teleoperators & virtual environments*, 15(1):1–15, 2006.

[16] L.-P. Cheng, E. Ofek, C. Holz, H. Benko, and A. D. Wilson. Sparse haptic proxy: Touch feedback in virtual environments using a general passive prop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 3718–3728, 2017.

[17] X. de Tinguy, C. Pacchierotti, M. Marchal, and A. Lécuyer. Toward universal tangible objects: Optimizing haptic pinching sensations in 3d interaction. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 321–330. IEEE, 2019.

[18] H. G. Debarba, R. Boulic, R. Salomon, O. Blanke, and B. Herbelin. Self-attribution of distorted reaching movements in immersive virtual reality. *Computers & Graphics*, 76:142–152, 2018.

[19] H. G. Debarba, J.-N. Khoury, S. Perrin, B. Herbelin, and R. Boulic. Perception of redirected pointing precision in immersive virtual reality. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 341–346. IEEE, 2018.

[20] B. Eckstein, E. Krapp, and B. Lugrin. Towards serious games and applications in smart substitutional reality. In *2018 10th International Conference on Virtual Worlds and Games for Serious Applications (VS-Games)*, pp. 1–8. IEEE, 2018.

[21] S. Esmaeili, B. Benda, and E. D. Ragan. Detection of scaled hand interactions in virtual reality: The effects of motion direction and task complexity. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 453–462. IEEE, 2020.

[22] G. T. Fechner. *Elemente der psychophysik*, vol. 2. Breitkopf u. Härtel, 1860.

[23] M. Feick, N. Kleer, A. Zenner, A. Tang, and A. Krüger. Visuo-haptic illusions for linear translation and stretching using physical proxies in virtual reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2021.

[24] B. R. Glowacki, R. Freire, L. M. Thomas, M. O'Connor, A. Jamieson-Binnie, and D. R. Glowacki. An open source etextile vr glove for real-time manipulation of molecular simulations. *arXiv preprint arXiv:1901.03532*, 2019.

[25] E. J. Gonzalez and S. Follmer. Investigating the detection of bimanual haptic retargeting in virtual reality. In *25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–5, 2019.

[26] T. Grechkin, J. Thomas, M. Azmandian, M. Bolas, and E. Suma. Revisiting detection thresholds for redirected walking: Combining translation and curvature gains. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 113–120, 2016.

[27] M. Guthold, M. R. Falvo, W. G. Matthews, S. Paulson, S. Washburn, D. A. Erie, R. Superfine, F. Brooks, and R. M. Taylor. Controlled manipulation of molecular samples with the nanomanipulator. *IEEE/ASME transactions on mechatronics*, 5(2):189–198, 2000.

[28] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 452–458, 1994.

[29] T. Ishiyama and T. Kitahara. A prototype of virtual drum performance system with a head-mounted display. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 990–991. IEEE, 2019.

[30] S. A. Klein. Measuring, estimating, and understanding the psychometric function: A commentary. *Perception & psychophysics*, 63(8):1421–1455, 2001.

[31] L. Kohli. Redirected touching: Warping space to remap passive haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 129–130. IEEE, 2010.

[32] Y. Konishi, N. Hanamitsu, B. Outram, K. Minamizawa, T. Mizuguchi, and A. Sato. Synesthesia suit: the full body immersive experience. In *ACM SIGGRAPH 2016 VR Village*, pp. 1–1. 2016.

[33] R. W. Lindeman, J. L. Sibert, and J. K. Hahn. Hand-held windows: towards effective 2d interaction in immersive virtual environments. In *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*, pp. 205–212. IEEE, 1999.

[34] R. W. Lindeman, J. L. Sibert, and J. K. Hahn. Towards usable vr: an empirical study of user interfaces for immersive virtual environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 64–71, 1999.

[35] T. H. Massie, J. K. Salisbury, et al. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, vol. 55, pp. 295–300. Chicago, IL, 1994.

[36] W. A. McNeely. Robotic graphics: a new approach to force feedback for virtual reality. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, pp. 336–341. IEEE, 1993.

[37] V. Mercado, M. Marchai, and A. Lécuyer. Design and evaluation of interaction techniques dedicated to integrate encountered-type haptic displays in virtual environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 230–238. IEEE, 2020.

[38] N. C. Nilsson, A. Zenner, and A. L. Simeone. Propping up virtual reality with haptic proxies. *IEEE Computer Graphics and Applications*, 41(5):104–112, 2021.

[39] S. Razzaque, D. Swapp, M. Slater, M. C. Whitton, and A. Steed. Redirected walking in place. In *EGVE*, vol. 2, pp. 123–130, 2002.

[40] N. Ryu, W. Lee, M. J. Kim, and A. Bianchi. Elastick: A handheld variable stiffness display for rendering dynamic haptic response of flexible object. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 1035–1045, 2020.

[41] S. Shaphiro and M. Wilk. An analysis of variance test for normality. *Biometrika*, 52(3):591–611, 1965.

[42] T. Stebbins and E. D. Ragan. Redirecting view rotation in immersive movies with washout filters. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 377–385. IEEE, 2019.

[43] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE transactions on visualization and computer graphics*, 16(1):17–27, 2009.

[44] P. L. Strandholt, O. A. Dogaru, N. C. Nilsson, R. Nordahl, and S. Serafin. Knock on wood: Combining redirected touching and physical props for tool-based interaction in virtual reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020.

[45] M. Turvey, G. Burton, E. L. Amazeen, M. Butwill, and C. Carello. Perceiving the width and height of a hand-held object by dynamic touch. *Journal of Experimental Psychology: Human Perception and Performance*, 24(1):35, 1998.

[46] K. Ushiyama, A. Takahashi, and H. Kajimoto. Modulation of a hand-held object's property through proprioceptive stimulation during active arm movement: Proprioceptive modulation of a hand-held object's property. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–6, 2021.

[47] R. J. van Beers, D. M. Wolpert, and P. Haggard. When feeling is more important than seeing in sensorimotor adaptation. *Current biology*, 12(10):834–837, 2002.

[48] D. Wang, K. Ohnishi, and W. Xu. Multimodal haptic display for virtual reality: A survey. *IEEE Transactions on Industrial Electronics*, 67(1):610–623, 2019.

[49] F. A. Wichmann and N. J. Hill. The psychometric function: I. fitting, sampling, and goodness of fit. *Perception & psychophysics*, 63(8):1293–1313, 2001.

[50] F. A. Wichmann, F. Jäkel, and J. Wixted. Methods in psychophysics. *Stevens handbook of experimental psychology and cognitive neuroscience*, 5:265–306, 2018.

[51] J. Yang, H. Horii, A. Thayer, and R. Ballagas. Vr grabbers: ungrounded haptic retargeting for precision grabbing tools. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pp. 889–899, 2018.

[52] A. Zenner, H. M. Kriegler, and A. Krüger. Hart-the virtual reality hand redirection toolkit. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2021.

[53] A. Zenner and A. Krüger. Shifty: A weight-shifting dynamic passive haptic proxy to enhance object perception in virtual reality. *IEEE transactions on visualization and computer graphics*, 23(4):1285–1294, 2017.

[54] A. Zenner and A. Krüger. Estimating detection thresholds for desktop-scale hand redirection in virtual reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 47–55. IEEE, 2019.

[55] A. Zenner, K. P. Reqitz, and A. Krüger. Blink-suppressed hand redirection. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 75–84. IEEE, 2021.

[56] A. Zenner, K. Ullmann, and A. Krüger. Combining dynamic passive haptics and haptic retargeting for enhanced haptic feedback in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2627–2637, 2021.

[57] Y. Zhao, C. L. Bennett, H. Benko, E. Cutrell, C. Holz, M. R. Morris, and M. Sinclair. Enabling people with visual impairments to navigate virtual reality with a haptic and auditory cane simulation. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pp. 1–14, 2018.